Network Dimensioning Based on Percentage of Access Link Capacity Carrying Premium Traffic

Vasilios A. Siris and Georgios I. Fotiadis

Foundation for Research and Technology - Hellas (FORTH), Institute of Computer Science and
University of Crete, Department of Computer Science
P.O. Box 1385, GR 711 10 Heraklion, Crete, Greece
Tel.: +30 2810 391726, fax: +30 2810 391601
{vsiris,fotiadis}@ics.forth.gr

***Abstract*— Network dimensioning is typically based on a traffic matrix that contains the bandwidth demand for each source-destination pair. We consider a different approach, which assumes that only the maximum amount of premium traffic that can be carried across each access link is known. Based on this information, we present three algorithms for determining the bandwidth required for premium traffic in the core network. The first algorithm intermediately estimates the traffic matrix based on the "gravity model". The other two algorithms compute the worst-case requirements for premium traffic when there are no link failures and fixed routing is used (second algorithm), and when there are link failures and alternative routing paths are available (third algorithm). An important feature of the two worst-case algorithms is that they allow admission control decisions to be based solely on the availability of bandwidth for premium traffic at the access links. We apply and compare the three algorithms for the Greek Research and Technology Network (GRNET), which currently uses the third algorithm for its premium IP service.**

I. INTRODUCTION

Quality of Service (QoS) has become a reality in most of today's packet switched networks. QoS defines the mechanisms used to provide bandwidth and delay/jitter guarantees to the subscribers of a multi-service communication network, which can simultaneously accommodate different types of traffic flows, such as voice-over-IP (VoIP), streaming video, and legacy best-effort Internet traffic (file transfer, e-mail and web traffic). The DiffServ model [1] is the most common approach for QoS provision. DiffServ introduces the notion of Per-Hop-Behavior (PHB) where each aggregate traffic flow is mapped to a traffic class, whose packets are treated equally. Each traffic class can represent a different service offered by the network provider. PHBs are implemented at routers using priority queuing and scheduling mechanisms.

Providing guarantees in large-scale DiffServ networks can be a complex task mainly due to the heterogeneity of the participating subscriber and provider networks, whose capacity can range from several Gbps to under one Mbps, as in the case of low-speed DSL links. This capacity mismatch can result in congestion at the network core and, more likely, at the network edge. One approach to address this problem is to limit the amount of premium traffic that can be carried over the access links, and implement an appropriate dimensioning and provisioning scheme for ensuring the guarantees for premium traffic; such an approach also avoids premium IP traffic from starving best-effort traffic. The maximum percentage of each access link that can be used for premium traffic can depend on the provider's service policy. Hence, in the case of a non-profit NREN (National Research and Educational Network), as is the case of the Greek Research and Technology Network (GRNET), subscribers with the same access link capacity can be allowed to use up to the same maximum percentage of the access link's capacity for carrying premium traffic, hence ensuring fairness in using the premium IP service. On the other hand, a provider selling network access services can offer different pricing options that correspond to a different percentage of premium IP traffic that can be sent across the access link.

In this paper we present and investigate three dimensioning algorithms that estimate the required premium IP bandwidth in the core network, taking as input the maximum percentage of each access link's capacity that can carry premium traffic. The first algorithm intermediately estimates the traffic matrix based on the so-called "gravity model" [2], [3]. The other two algorithms compute the worst-case bandwidth requirements when there are no link failures and fixed routing is applied (second algorithm), and when there are link failures and alternative routing paths are available (third algorithm). An important advantage of the latter two algorithms is that admission control decisions can be based solely on the availability of bandwidth for premium traffic at the access links. This greatly simplifies the core network, which only requires to give higher priority to premium traffic, through the provision of appropriate PHB mechanisms. Additionally, the three algorithms are independent of the particular routing algorithm implemented in the core network: the first and second algorithm assumes that the routing paths are known, whereas the third algorithm computes the worst-case bandwidth requirements when alternate routing paths can bypass failed links. The third algorithm is currently used for GRNET's premium IP service [4], where requests for premium connections are made through a web-based tool that takes admission control decisions based solely on the access

---

This work has been supported in part by the Greek Research and Technology Network (GRNET).

link capacities, their maximum percentage that can be used for premium traffic, and the active premium IP connections.

The remainder of this paper is organized as follows. Section II identifies related work in the field of network dimensioning and capacity planning. Section III provides the network definitions and the provisioning mechanisms used, while Section IV presents the three dimensioning algorithms we investigate. Section V presents the results from the application of the algorithms to the GRNET[1] network and, finally, Section VI concludes the paper.

II. RELATED WORK

Several methods have been proposed in the literature for estimating the traffic demand over a network's backbone in order to efficiently dimension network capacity. The traffic demand is typically expressed using a two-dimension traffic matrix that quantifies the demand between all source and destination pairs in the network; e.g. see [3], [5] and the references therein. The computation requires measurements of the actual traffic, e.g. at the ingress and egress interfaces. Rather than considering the point-to-point demand, the work of [6] considers the traffic demand originating from an ingress router to a set of egress routers; such a model is appropriate for transit domains, where a destination is reachable from multiple egress routers. The above approaches cannot be applied to premium IP services, which have not been offered before, hence for which no usage data is available.

One approach for network dimensioning of premium IP services involves the formulation of an optimization problem involving link costs. However, consideration of delay and loss constraints lead to an NP-complete problem, whose solution requires various heuristic assumptions [7]. Optimization objectives involving link costs can also be considered when premium IP traffic co-exists with best-effort traffic [8], [9].

Unlike the above works, the dimensioning approach we consider does not require knowledge of a traffic matrix between source-destination pairs, but only the maximum percentage of the access link capacity that can be used for premium IP traffic; moreover, the second and third algorithm we present find the worst-case bandwidth requirements, independent of the underlying routing algorithm. On the other hand, the works of [7], [8], [9] consider the capacity dimensioning problem in conjunction with routing. Such joint consideration would lead to more optimal dimensioning; however, in current networks that typically over-dimension some of the core network links carrying both best-effort and premium traffic, the added complexity required to modify the routing algorithm is likely to outweigh the increased efficiency.

III. DEFINITIONS AND MECHANISMS

In this section we provide the definitions used in the remaining sections, and describe the provisioning mechanisms used at the core and the edge of the network to meet the guarantees for premium traffic.

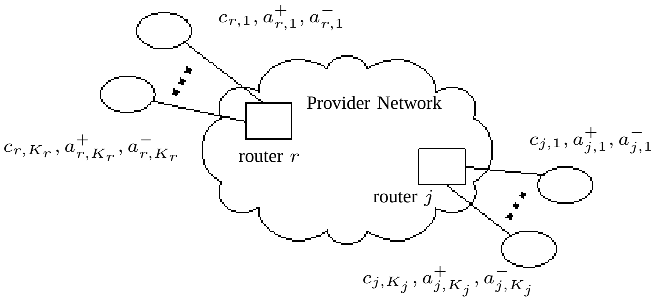

Fig. 1. A provider's edge router r is connected to K_r subscriber networks. The k'th subscriber network connected to router r has capacity $c_{r,k}$. A maximum percentage $a^{+}_{r,k}, a^{-}_{r,k}$ of this access link capacity can carry premium IP traffic in the ingress, egress direction, respectively.

A. Definitions

Let $G = (R, L)$ be a directed graph representing the network topology, where R is the set of routers and L is the set of links in the network interconnecting the routers in R. The capacity of link l is C_l.

Every router $r \in R$ of the provider network has a number K_r of subscriber networks connected to it, Figure 1. Each subscriber network is connected to the provider's edge router through an access link with capacity $c_{r,k}$, where $r \in R$ and $1 \leq k \leq K_r$. A maximum percentage of this access link capacity can be used to carry premium traffic. Let this percentage be $a^{+}_{r,k}$ in the ingress direction (from the subscriber to the provider) and $a^{-}_{r,k}$ in the egress direction (from the provider to the subscriber).

For each edge router r we define $I(r)$ as the maximum aggregate rate of premium traffic that can be inserted to the provider's network through router r from all its subscriber networks, and $O(r)$ as the maximum aggregate rate of premium traffic that can be ejected from the provider's network through router r to all its subscriber networks. Hence,

$$I(r) = \sum_{k=1}^{K_r} a^{+}_{r,k} \cdot c_{r,k} \; , \quad O(r) = \sum_{k=1}^{K_r} a^{-}_{r,k} \cdot c_{r,k} \; . \tag{1}$$

If the network does not contain any servers that only produce or consume premium IP traffic, we will have $a^{+}_{r,k} = a^{-}_{r,k}$, hence $I(r) = O(r), \forall r \in R$. Note, however, that the algorithms presented in Section IV do not require this assumption.

B. Provisioning mechanisms

In order to satisfy the requirements for premium IP services appropriate mechanisms have to be implemented both at the edge and the core of the network. These mechanisms include admission control, policing, and queueing.

The objective of admission control is to ensure that requests for premium IP connections are accepted only if there is available bandwidth. Note that admission control is applied in addition to defining a maximum percentage of each access link's capacity that can be used for premium traffic. Hence, it is not enough to check that the access link of a subscriber requesting a premium connection has enough bandwidth to accept the request. With the application of the second and

---

[1] The Greek Research and Technology Network (GRNET), `http://www.grnet.gr`

third dimensioning algorithms, which estimate the worst-case bandwidth requirements when there are no link failures and fixed routing is applied (second algorithm), and when there are link failures and alternative routing paths are available (third algorithm), the admission control decision needs to involve only the ingress and egress access links that the requested premium IP connection will traverse.

The second provisioning mechanism is policing, whose objective is to ensure that the traffic entering the network conforms to the maximum amount of traffic allowed for the particular connection. Policing is applied solely at the ingress router of the provider network, and is based on the source and destination subscriber network of the premium IP connection.

The third provisioning mechanism involves queueing, whose objective is to give higher priority to premium IP traffic. In a DiffServ network this is achieved through the Expedited Forwarding (EF) PHB mechanism. Recall that starvation of best-effort traffic is avoided by limiting the amount of premium IP traffic at the access links.

IV. NETWORK DIMENSIONING ALGORITHMS

In this section we present three dimensioning algorithms whose input is the maximum amount of premium traffic that can be carried across each access link.

A. Gravity-based Traffic Matrix (GTM) Dimensioning Algorithm

The Gravity-based Traffic Matrix (GTM) dimensioning algorithm uses the maximum amount of premium traffic carried across each access link to estimate a traffic matrix using the so-called "gravity" model. The gravity model originated in social science, but has been successfully applied in telephony to accurately model the telephone traffic exchanged between two cities based on their number of area codes, and in IP networks to estimate the traffic matrix from network measurements; e.g. see [2], [3], [5] and the references therein. According to this model the amount of premium traffic $M(s,d)$ originating from router s and destined to router d is proportional to the amount of premium traffic that can be accepted by router d, hence

$$M(s,d) = I(s) \cdot \frac{O(d)}{\sum_{x \in R} O(x)} .$$

If we define $P(s,d)$ to be the path from s to d, then the bandwidth requirement for premium IP traffic $B(l)$ at each link $l \in L$ can be computed using the following Algorithm 1:

Algorithm 1 Gravity-based Traffic Matrix (GTM) Algorithm

```
for all $s \in R$ do
  for all $d \in R$ do
    for all $l \in P(s,d)$ do
      $B(l) += M(s,d)$
    end for
  end for
end for
```

Algorithm 1 adds to each link belonging to the path from router s to d the amount of traffic $M(s,d)$ estimated using the gravity model assumption. In Section V where we apply the proposed algorithms to the GRNET network, the GTM algorithm will be used as a baseline comparison since it represents an accurate estimate of the bandwidth requirements for premium IP traffic, if the gravity model's proportionality assumption used for estimating the traffic matrix is accurate.

B. Worst-Case (WC) Dimensioning Algorithm with no Link Failures and Fixed Routing

This worst-case (WC) dimensioning algorithm defines an upper bound for the bandwidth required for premium traffic on all network links, when no link failures occur and fixed routing is assumed. The algorithm's input are variables I and O, defined in (1). The algorithm assumes the path $P(s,d)$ between each source-destination pair is static and known.

The WC algorithm relies on the calculation of two variables: $B_I(l)$ denotes the aggregate traffic that can be injected from all upstream routers to link l (recall that links are directed), and is given by

$$B_I(l) = \sum_{s:\exists d \in R: l \in P(s,d)} I(s) .$$

Similarly, $B_O(l)$ denotes the aggregate traffic that can be ejected to routers that are downstream of link l. So:

$$B_O(l) = \sum_{d:\exists s \in R: l \in P(s,d)} O(d) .$$

The bandwidth allocations for every link is determined through the following lemma.

Lemma 1: The maximum bandwidth requirements for premium IP traffic on link $l \in L$ is given by the minimum of $B_I(l)$ and $B_O(l)$.

Proof: Suppose that $\exists l \in L$ with $B_I(l) = a$ and $B_O(l) = b$, and $a > b$. It is not possible for link l to carry premium traffic with rate a, even though upstream routers can inject such an amount traffic, because downstream routers can only accept traffic with maximum rate $b < a$. Hence, the maximum amount of premium traffic that can flow through link l is b. Similarly, if $a < b$ then the maximum premium traffic is limited by the amount of traffic that can be injected in the link by its upstream routers, which is a.

The pseudocode for the worst-case (WC) dimensioning algorithm is shown in Algorithm 2.

C. Worst-Case Dimensioning Algorithm with Link Failures (WC-LF)

Next we present a worst-case dimensioning algorithm that considers link failures and assumes alternate routes exist; this is the case when the network contains loops, i.e. a set of links whose endpoints form a closed loop.

Let T be a spanning tree of the directed graph representing the network topology; T contains a single path from any source router to any destination router. The algorithm presented in the previous subsection can be used to calculate the maximum bandwidth requirements $B_T(l)$ for each link l belonging to the spanning tree T. The overall worst-case

Algorithm 2 Worst-Case (WC) Dimensioning Algorithm

```
for all $s \in R$ do
  for all $l \in L$ do
    if $\exists d \in R$ such that $l$ in $P(s,d)$ then
      $B_I(l) += I(s)$
    end if
  end for
end for

for all $d \in R$ do
  for all $l \in L$ do
    if $\exists s \in R$ such that $l$ in $P(s,d)$ then
      $B_O(l) += O(d)$
    end if
  end for
end for

for all $l \in L$ do
  $B(l) = \min\{B_I(l), B_O(l)\}$
end for
```

bandwidth requirement $B^*(l)$ for link l can be found by taking the maximum of $B_T(l)$ over all spanning trees T; this considers all possible link failures and alternate routes. The correctness of this approach for finding the worst-case bandwidth requirements follows directly from Lemma 1.

Let V be the set of loops, and $V^* = \{l : l \in v, v \in V\}$ the set of all links belonging to loops. Next we propose an algorithm for calculating the worst-case bandwidth requirements on all links. We will initially assume that the loops in the network are disjoint, i.e. they do not share common links. We consider this special case since many wide-area networks, such as GRNET, have such a topology. Later we will present a more general algorithm for the general case.

In the simple case of disjoint loops, all the different spanning trees can be found by sequentially marking each link in a loop as down, and applying the worst-case Algorithm 2. The resulting pseudocode is shown in Algorithm 3. At the end of the algorithm, the maximum bandwidth requirement for link l is given by $B^*(l)$.

Algorithm 3 Worst-Case Dimensioning Algorithm with Link Failures (WC-LF)

```
while $\exists l \in V^* : l$ not processed do
  for all $v \in V$ do
    if $\exists l \in v : l$ not processed then
      mark $l$ as down
      mark $l$ as processed
    end if
  end for
  run $WC$ Algorithm 2
  for all $l \in L$ do
    $B^*(l) = \max\{B(l), B^*(l)\}$
  end for
  mark all links as up
end while
```

D. *WC-LF+ Algorithm*

Next we extend Algorithm 3 to the case of network topologies that contain loops which share common links. The approach followed by this algorithm extends the WC-LF algorithm by marking more than one links in the original graph as down, such that the remaining links create a single path between the different router pairs.

We define $P^*(s,d) = \{P(s,d)\}$ to be all the possible paths from router s to d, and as before $V^* = \{l : l \in v, v \in V\}$ the set of all links belonging to loops. The pseudocode for the WC-LF+ algorithm is shown below. At the end of the algorithm, the maximum bandwidth requirement for link l is given by $B^*(l)$.

Algorithm 4 WC-LF+ Dimensioning Algorithm

```
for all $s \in R$ do
  for all $d \in R$ do
    for all $p \in P^*(s,d)$ do
      mark one link $l$ from every $p' \in P^*(s,d) - \{p\}$
      as down, s.t. $l \in V^*$ and $l$ belongs to a single loop
      run $WC$ Algorithm 2
      for all $l \in L$ do
        $B^*(l) = \max\{B(l), B^*(l)\}$
      end for
    end for
  end for
end for
```

It is worth noting that the two worst-case algorithms are equivalent to finding for each link l, the worst-case bandwidth requirements over all cuts of the network graph that divide it into two disjoint subgraphs, and the two end-points of link l are contained in different subgraphs [4].

V. APPLICATION TO THE GREEK RESEARCH AND TECHNOLOGY NETWORK (GRNET)

In this section we present and discuss results from the application of the three dimensioning algorithms described in the previous section to the GRNET network, whose premium IP service is used for VoIP and other delay-sensitive traffic. The Gravity-based Traffic Matrix (GTM) algorithm yields an accurate estimate of the bandwidth requirements for premium traffic, if the gravity model's proportionality assumption used for estimating the traffic matrix is accurate. Hence, comparison of the GTM algorithm with the two worst-case dimensioning algorithms (Algorithm 2 - WC, and Algorithm 3 - WC-LF) will show when and by how much these two algorithms overestimate the required bandwidth for premium traffic. Note that since the GRNET network does not contain loops with common links, we apply Algorithm 3 for estimating the worst-case premium IP bandwidth requirements in the case of link failures. Also, because loops do not share common links, there are two routes between routers belonging to loops.

The GRNET network is shown in Figure 2. The rectangular nodes represent core routers with high speed interfaces (up to 2.5 Gbps) that typically contain a large number of subscriber networks, while the round nodes are routers with lower speed

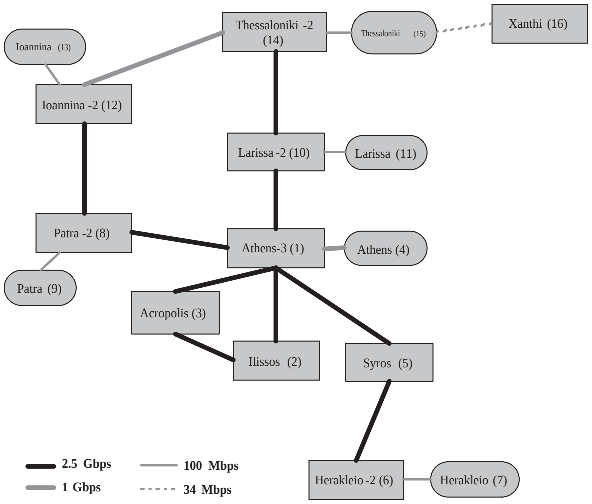

Fig. 2. The GRNET Network

interfaces (up to 100 Mbps, except the Athens router that has a 1 Gbps interface). The access links of subscriber networks are highly heterogenous, ranging from speeds of 1 Gbps (University of Athens, connected to the *Acropolis* router) to 2 Mbps (Institute of Molecular Biology Crete - IMBC, connected to the *Herakleio* router). The maximum percentage of the access link's capacity that can be used to carry premium traffic is shown in Table I; we assume that the percentage is the same in the ingress and egress direction.

Table II shows the subscriber networks of router *Herakleio* and their corresponding access link capacities. The aggregate amount of traffic that can be injected in the network, I, and ejected from the network, O, from the *Herakleio* router can be calculated using (1) and Table II. Hence,

$$I(\textit{Herakleio}) = O(\textit{Herakleio}) =$$

$$0.2 \cdot 34 + 0.2 \cdot 34 + 0.2 \cdot 4 + 0.2 \cdot 2 = 14.8 \text{ Mbps}\,.$$

The values of the variables I and O for the other routers in GRNET are shown in Table III; for all routers, I is equal to O since we have assumed that the percentage of the access link that can be used for premium traffic is the same in the ingress and egress direction. Table III also shows each router's ID that appears in the notation 'link $a - b$', which represents the link connecting routers with ID a and b.

TABLE I

ACCESS LINK PERCENTAGES FOR PREMIUM IP TRAFFIC

Access Link Capacity c	Percentage Reserved
1 Gbps	1%
100 Mbps	10%
$\leq$ 34 Mbps	20%

TABLE II

SUBSCRIBER NETWORKS AT NODE *Herakleio*

Subscriber Network	Access Link Capacity (Mbps)
UoC	34
FORTH	34
TEIHER	4
IMBC	2

TABLE III

GRNET ROUTERS

Router Name	Router ID	I, O (Mbps)
Athens-3	1	50.00
Ilissos	2	50.00
Acropolis	3	50.00
Athens	4	12.54
Syros	5	20.00
Herakleio-2	6	20.00
Herakleio	7	14.80
Patra-2	8	11.20
Patra	9	10.20
Larissa-2	10	10.00
Larissa	11	7.20
Ioannina-2	12	10.00
Ioannina	13	7.30
Thessaloniki-2	14	20.00
Thessaloniki	15	17.00
Xanthi	16	24.80

TABLE IV

PREMIUM IP BANDWIDTH REQUIREMENTS

Link ID	GTM	(%)	WC	(%)	WC-LF	(%)
1-2	35.08	1.40	50.00	2.00	100.00	4.00
1-3	35.08	1.40	50.00	2.00	100.00	4.00
1-4	12.07	1.21	12.54	1.25	12.54	1.25
1-5	45.84	1.83	54.80	2.19	54.80	2.19
1-8	26.20	1.05	38.70	1.55	117.70	4.71
1-10	52.35	2.09	79.00	3.10	117.70	4.71
2-3	7.46	0.30	50.00	2.00	50.00	2.00
5-6	31.19	1.25	34.80	1.39	34.80	1.39
6-7	14.15	14.2	14.80	14.8	14.80	14.8
8-9	9.89	9.89	10.20	10.2	10.20	10.2
8-12	16.27	0.65	79.10	3.16	96.30	3.85
10-11	7.05	7.05	7.20	7.20	7.20	7.20
10-14	44.15	1.77	79.10	3.16	100.50	4.02
12-13	7.14	7.14	7.30	7.30	7.30	7.30
12-14	8.03	0.80	38.70	3.87	79.00	7.90
14-15	36.58	36.6	41.80	41.8	41.80	41.8
15-16	22.96	67.5	24.80	72.9	24.80	72.9

Results from the application of the three dimensioning algorithms appear in Table IV, which shows the maximum amount of bandwidth and the corresponding percentage of each link's capacity required for premium traffic that is estimated by the three dimensioning algorithms. Because the percentage of the access link capacity used for premium traffic is the same in the ingress and egress direction, the bandwidth requirements are symmetric, hence the maximum premium IP bandwidth for link $a - b$ is equal to that of link $b - a$.

As expected, the bandwidth requirements for premium IP traffic estimated by the Gravity-based Traffic Matrix (GTM) algorithm are smaller than the two worst-case dimensioning algorithms WC and WC-LF. Additionally, the bandwidth estimates for the WC-LF algorithm are equal to or higher than the other two algorithms; this is expected since the WC-LF algorithm considers alternate routes to bypass link failures, whereas the other two algorithms assume that links do not fail and consider fixed routes.

The percentage by which the two worst-case algorithms overestimate, compared to the GTM algorithm, the premium IP traffic bandwidth requirements is not the same for all links. In particular, the overestimation for links that do not belong to a network loop is up to approximately 14% (for link $14-15$).

Additionally, for these links the estimate by the two worst-case dimensioning algorithms is the same; this occurs because for links that do not belong to loops there are no alternate routes. Finally, note that except for links $14-15$ and $15-16$, the estimated percentage of premium IP traffic on all links not belonging to a loop are under 15%, which is less than the maximum percentage of 25% premium traffic typically suggested for supporting low loss and jitter [10]. On the other hand, for links $14-15$ and $15-16$ the estimated bandwidth for premium IP traffic is above 35% for all three dimensioning algorithms; this suggests that the capacity for these two links should be increased in order to effectively accommodate the full demand for premium IP services.

For the links that belong to loops, the overestimation for premium IP bandwidth by the two worst-case algorithms is higher. However, the overestimation is not the same for all links. Hence, the overestimation of the WC algorithm for links $1-10$, $10-14$, and $1-8$ is lower (up to 80% for link $10-14$) than for links $8-12$ and $12-14$ (up to 380% for link $8-12$). This is attributed to the fact that routers 8 and 12 contain a small aggregate subscriber rate, and the premium IP bandwidth estimates of the GTM algorithm for links $8-12$ and $12-14$ are small because these links are not contained in many routing paths, as is the case of links $1-10$, $10-14$, and $1-8$. The latter observation also explains why the overestimation for link $2-3$ (570%), which is contained only in the path from router 2 to 3, is larger than for links $1-2$ and $2-3$ (42%). Also, note that for these links the WC-LF algorithm overestimates the premium IP bandwidth compared to the WC algorithm by 100%; this is to be expected since, unlike the WC algorithm, the WC-LF algorithm considers the case of link failures and there are two routes between routers belonging to a loop. It is important to note that for all the above links the estimated premium IP bandwidth by all three algorithms is less than 8%, which is less than the percentage 25% mentioned above for supporting low loss and jitter.

From the above results we can conclude that the maximum percentage of the access link capacities for premium IP traffic shown in Table I can be supported on the core links of GRNET's network, except for links $14-15$ and $15-16$ where the percentage required for premium IP traffic even with the Gravity-based Traffic Matrix (GTM) algorithm is above 35%, which is rather too high; this result suggests that the capacity for these two links should be increased. The discussion in this section shows how the proposed dimensioning algorithms can be used in practise, to determine if the capacity of core links in a wide-area network needs to be increased, and by how much, or to estimate the minimum amount of capacity that is available for best-effort traffic with and without link failures.

VI. CONCLUDING REMARKS

We presented three algorithms for dimensioning a network supporting premium IP traffic, when only the maximum amount of premium traffic that can be carried across each access link is known. The Gravity-based Traffic Matrix (GTM) algorithm is expected to be accurate if the proportionality assumption of the gravity model is accurate; other works have shown that it is quite accurate for telephony and best-effort traffic. The other two dimensioning algorithms estimate the worst-case bandwidth requirements, hence can be more appropriate when premium services are first rolled out and enable a provider to offer a higher level of assurances for premium IP traffic compared to the GTM algorithm. Moreover, they allow admission control decisions to be based solely on the availability of bandwidth at the access links. Which of the two worst-case dimensioning algorithms should be used will depend on the probability of link failures, since only the WC-LF (Worst-Case with Link Failures) dimensioning algorithm considers link failures, hence its bandwidth estimates can support guarantees in such cases. We compared the estimates of the three algorithms through their application to an actual wide-area network, the Greek Research and Technology Network (GRNET). The comparison gives intuition of when the worst-case algorithms overestimate the required premium IP bandwidth, and demonstrate how the algorithms can be used to determine if the capacity in core network links needs to be increased, and for estimating the minimum bandwidth for best-effort traffic with and without link failures.

The worst-case dimensioning algorithms assumed that traffic from a particular source router can be destined to any destination router. It would be interesting to adapt the algorithms in this paper to the case where the traffic from a source router can be destined to a set of destination routers, rather than to all routers; such a model is appropriate for transit networks. Another area for further work is the application of the proposed dimensioning algorithms to more complex network topologies, such as the European-wide GEANT network[2].

REFERENCES

[1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Services," *RFC 2475*, December 1998.

[2] D. Lam, D. Cox, and J. Widom, "Teletraffic Modeling for Personal Communications Services," *IEEE Commun.*, vol. 35, no. 2, Feb. 1997.

[3] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg, "Fast accurate computation of large-scale IP traffic matrices from link loads," in *Proc. of ACM SIGMETRICS'03*.

[4] A. Varvitsiotis, V. A. Siris, D. N. Primpas, G. I. Fotiadis, A. C. Liakopoulos, and C. J. Bouras, "Techniques for DiffServ-based QoS in Hierarchically Federated MAN Networks: the GRNET Case," in *Proc. of IEEE LANMAN 2005*.

[5] A. Gunnar, M. Johansson, and T. Telkamp, "Traffic Matrix Estimation on a Large IP Backbone - A Comparison on Real Data," in *Proc. of Internet Measurement Conference 2004*, October 2004.

[6] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving Traffic Demand for Operational IP Networks: Methodology and Experience," in *Proc. of ACM SIGCOMM'00*.

[7] P. Trimintzios, T. Bauge, G. Pavlou, L. Georgiadis, R. Egan, and P. Flegkas, "Quality of Service Provisioning for Supporting Premium Services in IP Networks," in *Proc. of IEEE Globecom 2002*.

[8] K. Wu and D. Reeves, "Capacity Planning of DiffServ Networks with Best-Effort and Expedited Forwarding Traffic," *Telecommunication Systems*, vol. 25, pp. 193–207, 2004.

[9] ——, "Link Dimensioning and LSP Optimization for MPLS Networks Supporting DiffServ EF and BE traffic classes," *18th International Teletraffic Congress*, 2003.

[10] C. Filsfils and J. Evans, "Engineering a multiservice IP backbone to support tight SLAs," *Computer Networks*, vol. 40, pp. 131–148, 2002.

---

[2] The European Academic and Research Backbone GEANT, `http://www.geant.net`